arXiv:1704.03009v1 [q-bio.QM] 10 Apr 2017

Deep neural networks automate detection for tracking of submicron scale particles in 2D and 3D

Jay M. Newby* Alison M. Schaefer† Phoebe T. Lee‡
M. Gregory Forest* Samuel K. Lai†

April 12, 2017

Abstract

Particle tracking is a powerful biophysical tool that requires conversion of large video files into position time series, i.e. traces of the species of interest for data analysis. Current tracking methods, based on a limited set of input parameters to identify bright objects, are ill-equipped to handle the spectrum of spatiotemporal heterogeneity and poor signal-to-noise ratios typically presented by submicron species in complex biological environments. Extensive user involvement is frequently necessary to optimize and execute tracking methods, which is not only inefficient but introduces user bias. To develop a fully automated tracking algorithm, we developed a convolutional neural network comprised of over 50,000 parameters, and employed deep learning to train the network on a diverse portfolio of video conditions. The neural network tracker, with no user-dependent input parameters, offered superior tracking performance, with exceptionally low lower false positive and false negative rates on both 2D and 3D simulated videos and 2D experimental videos of difficult-to-track species.

Introduction

In particle tracking experiments, high-fidelity tracking of an ensemble of species recorded by high-resolution video microscopy can reveal critical information about species transport within cells or mechanical and structural properties of the surrounding environment. For instance, particle tracking has been extensively used to measure the real-time penetration of pathogens across physiological barriers [1, 2, 3], to facilitate the development of nanoparticle systems for transmucosal drug delivery [4, 5], to explore dynamics and organization of domains of chromosomal DNA in the nucleus of living cells [6], and to characterize the micro- and meso- scale rheology of complex fluids via engineered probes [7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. The extraction of individual traces from raw videos is generally divided into two steps: (i) identifying the precise locations of particle centers from each frame of the video, and (ii) linking these particle centers across sequential frames into tracks or paths. There has been significant progress towards the goal of fully automated tracking, and dozens of methods are currently available that can automatically process videos, given a predefined set of adjustable parameters [18, 19]. Nonetheless, when presented with videos containing spatiotemporal heterogeneity (see Fig. 1) such as variable background intensity, photobleaching or low signal-to-noise ratio (SNR), the set of parameters used by a given method must be optimized for each set of video conditions, or even each video, which is highly subjective in the absence of ground truth.

Parameter optimization is time consuming and requires substantial user guidance. Furthermore, when applied to experimental videos, user input is still frequently needed to remove

*Department of Mathematics and Applied Physical Sciences, University of North Carolina–Chapel Hill, Chapel Hill, NC 27599
†Division of Pharmacoengineering and Molecular Pharmaceutics, Eshelman School of Pharmacy, University of North Carolina–Chapel Hill, Chapel Hill, NC 27599
‡UNC-NCSU Joint Department of Biomedical Engineering, University of North Carolina–Chapel Hill, Chapel Hill, NC 27599

phantom traces (false positives) or add missing traces (false negatives) (Fig. 2A-B). Thus, instead of providing full automation, current software is perhaps better characterized as facilitating supervised particle tracking, requiring substantial human interaction that is time consuming and costly. More importantly, the results can be highly variable, even for the same input video (Fig. 2C-E).

During the 2012 IEEE International Symposium on Biomedical Imaging, an open competition was held to develop and compare different particle tracking techniques [19]. One of the main outcomes of the study was the recommended process for optimizing particle tracking methods to experimental data. Specifically, the authors proposed to first generate simulated videos that match as closely as possible to the observed experimental conditions. Second, a tracking method suitable for those conditions is selected and applied to the simulated videos, and the error is quantitatively assessed. Third, the parameters in the tracking method are optimized to minimize the tracking error. Finally, once the parameters have been optimized to minimize tracking error on simulated data, the same parameters are used to analyze experimental videos.

To overcome the need to optimize for each video condition, we take the aforementioned methodology to the next logical step: instead of optimizing a tracking algorithm for a specific microscopy conditions, we compile a large portfolio of simulations that encompasses the full spectrum of potential variations that can be encountered in particle tracking experiments. Existing methods are designed with as few parameters as possible to make the software simple to use for humans, and a single set of parameters can usually be found for a specific microscopy conditions (SNR, size, shape, etc.) that identifies objects of interest. Nevertheless, a limited parameter space compromises the ability to optimize the method for an entire portfolio. An alternative approach is to construct an algorithm with thousands of parameters, and employ machine learning to optimize the algorithm to perform well under all conditions represented in the portfolio.

Here, we introduce an artificial neural network algorithm–called a convolutional neural network (CNN)– comprised of a 3-layer architecture and over 50,000 tunable parameters. A CNN is a type of feed-forward artificial neural network designed to process information in a layered network of connections. In recent years, CNNs have become the state-of-the-art for object recognition in computer vision [20, 21]. All of the neural network's tunable parameters are optimized using machine learning, which means there are never any parameters that the user needs to adjust for particle localization. The result is a highly optimized network that can perform under a wide range of conditions without any user supervision. To demonstrate accuracy, we test the neural network tracker on a large set of challenging videos that span a wide range of conditions, including variable background, particle motion, particle size, and low SNR

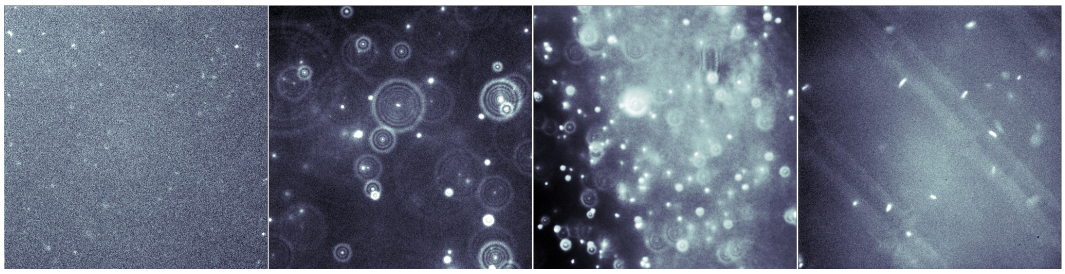

Figure 1: **Sample frames from experimental videos, highlighting some of the challenging conditions for particle tracking.** (from left to right) 50nm particles captured at low SNR, 200nm particles with diffraction disc patterns, variable background intensity, and ellipsoid PSF shapes from $1-2\mu m$ Salmonella.

Simulation of 4D greyscale image data

To train the network on a wide range of video conditions, we developed new video simulation software that accounts for a large range of conditions found in particle tracking videos (see

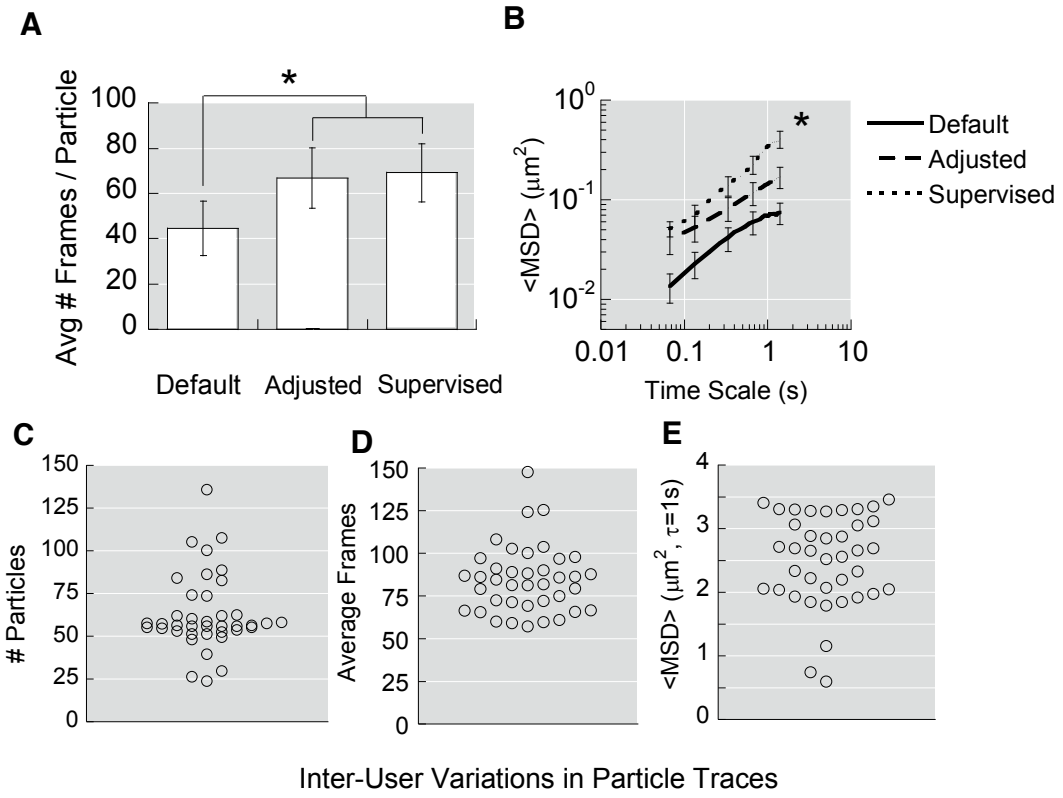

Figure 2: **The need for supervision in particle tracking, and inter-user variations in supervised tracking data.** Data represents the average of 4 movies of muco-inert 200 nm PEGylated polystyrene beads in human cervicovaginal mucus. Data from human supervised tracking (Supervised), which includes manually inspecting paths to remove false positives and minimize false negatives, is compared to results generated under default conditions of the tracking software (Default) and conditions manually adjusted by the user (Adj) to improve tracking accuracy. (A) average frames per particle; (B) ensemble-averaged geometric mean square displacements ($\langle$MSD$\rangle$) vs. time scale. Error bars represent standard error of the mean. * indicates statistically significant difference compared to 'Standard' ($p < 0.05$). (C-E) Inter-user variations in particle tracking. Different tracking software users were asked to analyze the same video of 200 nm bead in human cervicovaginal mucus. (A) Total particles tracked; (B) average frames tracked per particle; and (C) ensemble-averaged geometric mean square displacements ($\langle$MSD$\rangle$) at a time scale (τ) of 1s.

Fig. 1). The primary advance is to include simulations of how particles moving in 3D appear in a 2D image slice captured by the camera.

A standard camera produces images that are typically single channel (grey scale), and the image data is collected into four dimensional (three space and one time dimension) arrays of 16 bit integers. The resolution in the (x, y) plane is dictated by the camera and can be in the megapixel range. The resolution in the z coordinate is much smaller since each slice requires a piezo-electric motor to move the lense relative to the sample. A good piezo-electric motor is capable of moving between slices within a few milliseconds, which means that there is a trade off between more z-slices and the over all frame rate. For particle tracking, a typical video includes 10-50 z-slices per volume. The length of the video refers to the number of time points, i.e., the number of volumes collected. Video length is often limited by photobleaching, which slowly lowers the SNR as the video progresses.

To simulate a particle tracking video, we must first specify how particles appear in an image. The pixel intensities captured by a microscope and camera resulting from a particle centered at a given position (x, y, z) is called a point spread function (PSF), denoted by $\psi_{ijk}(x, y, z)$, where i, j, k are the pixel indices. The PSF becomes dimmer and less focused as the particle moves away from the plane of focus ($z = 0$). Away from the plane of focus, the

PSF also develops disc patterns caused by diffraction, which can be worsened by spherical aberration. While deconvolution can mitigate the disc patterns appearing in the PSF, the precise shape of the PSF must be known or unpredictable artifacts may be introduced into the image.

The shape of the PSF depends on several parameters that vary depending on the microscope and camera, including emitted light wavelength, numerical aperture, pixel size, and the separation between z-axis slices. It is not practical for the purpose of automatic particle tracking to develop a physical model based on optical physics that exposes these parameters. In practice, there are many additional factors that affect the PSF, such as the refractive index of the glass slide, of the lens oil (if oil-immersion objective is used), and of the medium containing the particles being imaged. The latter presents the greatest difficulty since biological specimens are often heterogeneous, and their optical properties are difficult to predict. The PSF can also be affected by particle velocity, depending on the duration of the exposure interval used by the camera. This makes machine learning particularly appealing, because we can simply randomize the shape of the PSF to cover a wide range of conditions, and the resulting CNN is capable of automatically 'deconvolving' PSFs without the need to know any of the aforementioned parameters.

Low SNR is an additional challenge for tracking of submicron size particles. High performance digital cameras are used to record images at a sufficiently high frame rate to resolve statistical features of particle motion. Particles with a hydrodynamic radius in the range of 10-100nm move quickly, requiring a small exposure time to minimize dynamic localization error (motion blur) [22]. Smaller particles also emit less light for the camera to collect. To train the neural network to perform in these conditions, we add Poisson shot noise with random intensity to the training videos. We also add slowly varying random background patterns (see Supplemental Figure 1).

An artificial neural network for particle localization

The 'neurons' of the artificial neural network are arranged in layers, which opperate on multi-dimensional arrays of data. Each layer output is 3 dimensional, with 2 spatial dimensions and an additional 'feature' dimension (see Fig. 3). Each feature within a layer is tuned to respond to specific patterns, and the ensemble of features are sampled as input to the next layer to form features that recognize more complex patterns. For example, the lowest layer is comprised of features that detect edges of varying orientation, and the second layer features are tuned to recognize curved lines and circular shapes (see Fig. 3).

Each neuron in the network processes information from spatially local inputs (either pixels of the input image or lower layer neurons). This enables a neuron to, figuratively speaking, see a local patch of the input image, which is smaller than the entire input image. The size of the image patch that affects the input to a given neuron is called its receptive field. The input and output, denoted by I_{ij} and O_{ij}, relationship for each neuron is given by

$$O_{ij} = F(\sum_{i',j'} w_{i',j'} I_{i+i',j+j'} - b), \tag{1}$$

where the kernel weights w_{ij} and output bias b are trainable parameters. Each layer has its own set of biases, one for each feature, and each feature has its own set of kernel weights, one for each feature in the layer directly below. The nonlinearity $F(\cdot)$ is a pre-specified function that determines the degree of 'activation' or output, we use $F(u) = \log(e^u + 1)$. Inserting nonlinearity in between each layer of neurons is necessary for CNNs to approximate nonlinear functions. The most common choice is called the rectified linear unit ($F(u \geq 0) = u$ and $F(u < 0) = 0$). Instead, we use a function with a similar shape that is also continuously differentiable, which helps minimize training iterations where the model is stuck in local minima [23].

The neural network is comprised of three layers; 12 features in layer one, 32 features in layer two, and the final two output features in layer three. The output of the neural net, denoted by q_{ijk}, can be interpreted as the probability of a particle centered at pixel (i, j, k). We refer to these as detection probabilities.

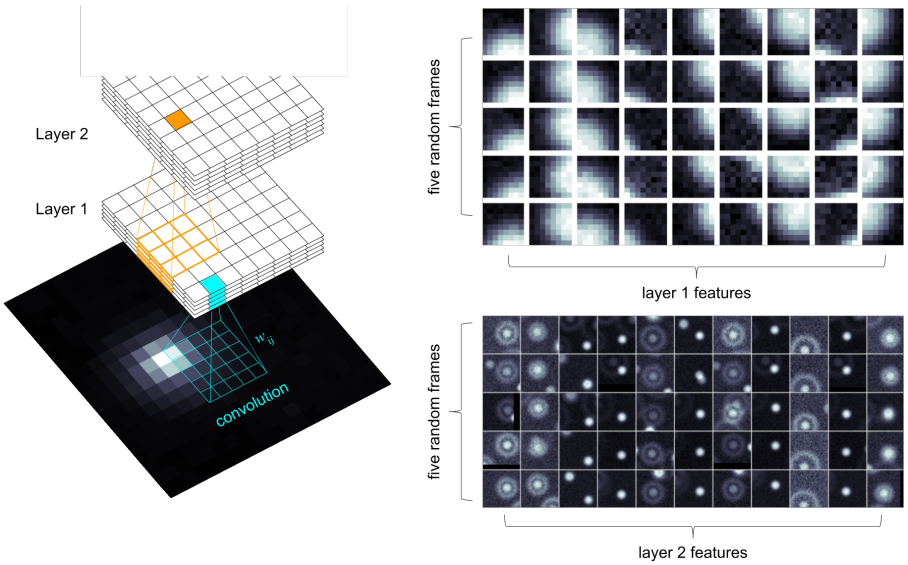

Figure 3: **The convolutional neural network.** (left) Diagram of the layered connectivity of the artificial neural network. (right) Features from layer one (**top**) and layer two (**bottom**). Five random frames from a video were processed through the neural net and image patches that maximally activated a given feature were recorded. Layer one features respond maximally to edges with varying orientation. Layer two features (only a subset shown) respond maximally to varying arrangements of circular spots surrounded or not surrounded by discs.

While it is possible to construct a network that takes 3D image data as input, it is not computationally efficient. Instead, the network is designed to process a single 2D image slice at a time (so that it can also be applied to the large set of existing 2D imaging data) while still maintaining the ability to perform 3D tracking. Constructing 3D output q_{ijk} is achieved by applying the network to each z-axis slice of the input image, the same way a microscope obtains 3D images by sequentially capturing each z-axis slice. Two or three dimensional paths can then be reconstructed from the network output as described below.

Advanced image classification tasks take color images as input. Most commonly, RGB images are used, which have three color channels: red, green, and blue. CNNs designed for image classification use all three channels as input to the first layer. Since microscopy videos are typically greyscale data, we replaced color channels with adjacent video frames. The neural net is comprised of three adjacent image frames so that each inference is obtained using information from the past and future to improve accuracy. Because detections include information from the past and future, the detection probabilities are reduced when a particle is not detected in the previous frame (the particle just appeared in the current frame) or is not detected in the following frame (the particle is about to leave the plane of focus). Below, we show how the detection probabilities can be used by the linking algorithm to improve its performance.

Optimizing the neural network parameters with machine learning

The values of the trainable parameters in the network, including the kernel weights and biases, are optimized through the process of learning. Using known physical models of particle motion and imaging, we simulate random particle paths and image frames that cover a wide range of conditions, including particle point spread function shape, variable background, particle number, particle mobility, and SNR. The 'ground truth' for each image consists of a binary image with pixels values $p_{ijk} = 1$ if $\|(j,i,k) - \mathbf{x}_n\| < 2$ and $p_{ijk} = 0$ otherwise. Each training image is processed by the neural net, and the corresponding output

is compared to the ground truth using the cross entropy error:

$$H[p,q] = -\frac{1}{N}\sum_{ijk}\left[p_{ijk}\log q_{ijk} + (1-p_{ijk})\log(1-q_{ijk})\right], \tag{2}$$

where N is the total number of pixels in the image. Cross entropy is (up to an additive constant that depends on p) a measure of how far the approximated distribution q is from the true distribution p. When $q = p$, the cross entropy reduces to the entropy of the true distribution p. Since p never changes for a given training video, our goal is to minimize $H[p,q]$ with respect to q over the entire training set of videos. At each iteration of the training procedure, a randomly generated training image is processed by the network, the error $H[p,q]$ is computed, and all of the trainable parameters are altered by a small amount (using the gradient decent method explained below) to reduce the observed error. This training procedure is repeated thousands of times until the error is minimized.

Suppose that all of the trainable parameters are arranged into the vector $\boldsymbol{\theta}$. The parameters are adjusted at the end of each training iteration t by computing the gradient of $\mathbf{g}_t = \nabla_\theta H[p_t, q_t]$. The gradient vector points in the direction of steepest rate of increase in the error, so the error can be reduced with $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - r\mathbf{g}_t$, where $r > 0$ is a predefined step size.

Generation of training images was performed in Python and training of the neural network was performed using Google's open source software package, Tensorflow. After training, the neural network is also deployed within Tensorflow, which executes the most computationally costly elements of the neural net tracker in highly optimized C++ code. Tensorflow can be easily adapted to use multiple cores of a CPU or GPU, depending on available hardware. On a standard desktop CPU with 4 cores, the neural net processes 512×512 resolution frames at ~ 5 frames per second.

Particle path linking

From the neural net output, we extract candidate particles along with their probabilities through thresholding the detection probabilities p_{ijk}, where ijk are the indices for each pixel of a single video frame. Detection probabilities for which $p_{ijk} < 0.5$ are eliminated. The remaining pixels contain local maxima ρ_n, which are extracted and identified as candidate particles using the method of connected components [24]. Connected sets of nearest neighbor pixels $\mathcal{P}_n$ above the threshold are collected as candidate particles. That is, $\mathcal{P}_n$ is a connected set and $p_{ijk} \geq 0.5$ for all $p_{ijk} \in \mathcal{P}_n$. Each candidate particle is assigned the largest pixel probability from its constituent pixel probabilities, i.e., $\rho_n = \max \mathcal{P}_n$. The position of each candidate particle is taken to be the center of mass given by,

$$\mathbf{x}_n = \frac{\sum_{p_{ijk}\in\mathcal{P}_n}(j,i,k)p_{ijk}}{\sum_{p_{ijk}\in\mathcal{P}_n}p_{ijk}}. \tag{3}$$

Note that there are alternative particle localization methods [25] that may increase accuracy. We have found that the center of mass method yields consistent sub-pixel accuracy of 0.6 pixels on average, which is sufficient for tracking tasks that require high accuracy such as micro-rheology. The next stage is to link candidate particles from one frame to the next.

The dynamics of particle motion can vary depending on the properties of the surrounding fluid and the presence of active forces (e.g., flagellar mediated swimming of bacteria and molecular motor cargo transport). In order to reconstruct accurate paths from a wide range of movement characteristics, we develop a minimal model that assumes only that particles move within a Gaussian range from one frame to the next. To accurately capture continuous motion sampled at discrete time intervals, dictated by the camera frame rate, the particle motion must be sufficiently small between image frames. Let $\mathcal{L}_t$ denote the set of linked particle pairs $(\mathbf{x}_t, \mathbf{x}_{t+1})$ together with their probabilities (ρ_t, ρ_{t+1}) in frame t to $t+1$. Let $\mathcal{N}_t^{\pm}$ be the set of probabilities for particles in frame t that are not linked to a particle in frame $t \pm 1$. Then, the log likelihood cost of the link assignments from frame t to frame $t+1$

is given by

$$
\begin{aligned}
L_t = &- \sum_{\mathbf{x}_t, \mathbf{x}_{t+1} \in \mathcal{L}_t} \frac{\|\mathbf{x}_t - \mathbf{x}_{t+1}\|^2}{2\sigma^2} \\
&+ \sum_{\rho_t, \rho_{t+1} \in \mathcal{L}_t} [\log \rho_t + \log \rho_{t+1}] \\
&- \sum_{\rho_t \in \mathcal{N}_t^+} \log(1 - \rho_t) - \sum_{\rho_{t+1} \in \mathcal{N}_{t+1}^-} \log(1 - \rho_{t+1}).
\end{aligned} \tag{4}
$$

The standard deviation σ is a user-specified parameter. Maximization of (4) can be formulated as a linear programming problem, which we solve using the Hungarian-Munkres algorithm [26].

Performance evaluation and comparison to existing software

We consider the primary goal for a high fidelity tracker to be accuracy (i.e., minimize false positives and localization error), followed by the secondary goal of maximizing data extraction (i.e., minimize false negatives and maximize path length). We therefore adopted error estimates focused on particle identification developed in [19]. To gauge accuracy, particle positions were matched to ground truth using optimal linear assignment. The algorithm finds the closest match between tracked and ground truth particle positions that are within a preset distance of 5 pixels; this is well above the sub-pixel error threshold of 1 pixel, but sufficiently small to ensure 1-1 matching. Tracked particles that did not match any ground truth particles were deemed false positives, and ground truth particles that did not match a tracked particle were deemed false negatives. To assess the performance of the neural net tracker, we analyzed the same videos using three different leading tracking software that are publicly available:

- **Mosaic (Mos)**: an ImageJ plug-in capable of automated tracking in 2D and 3D [27]
- **Icy**: an open source bio-imaging platform with pre-installed plugins capable of automated tracking in 2D and 3D [28, 29]
- **Video Spot Tracker (VST)**: a stand-alone application developed by the Center for Computer-Integrated Systems for Microscopy and Manipulation at UNC-CH capable of 2D particle tracking. VST also has a convenient graphic user interface that allows a user to add or eliminate paths (because human-assisted tracking is time consuming, 100 2D videos were randomly selected from the 500 video set)

Performance on simulated 2D videos

Because manual tracking by humans is subjective, our first standard for evaluating the performance of the neural net tracker (NN) and other publicly available software is to follow the recommended protocol set out in [19] and test on simulated videos, for which the ground truth particle paths are known. The test included 500 2D videos and 50 3D videos, generated using the video simulation methodology described in Section . Each 2D video contained 100 simulated particle paths for 50 frames at 512x512 resolution, (see Supplemental Figure 1). Each 3D video contained 20 evenly spaced z axis image slices of a 512x512x120 pixel region containing 300 particles. The conditions for each video were randomized, including variable background intensity, PSF radius (called particle radius for convenience), diffusivity, and SNR.

To assess the robustness of each tracking method/software, we used the same set of tracker parameters for all videos (see Supplementary material for further details). Scatter plots of the 2D test video results for neural network tracker, Mosaic, and Icy are shown in Fig. 4. For Mosaic, the false positive rate was generally quite low ($\sim$2%) when SNR > 3, but showed a marked increase to $>$20% for SNR < 3 (Fig. 4A). The average false negative rates were

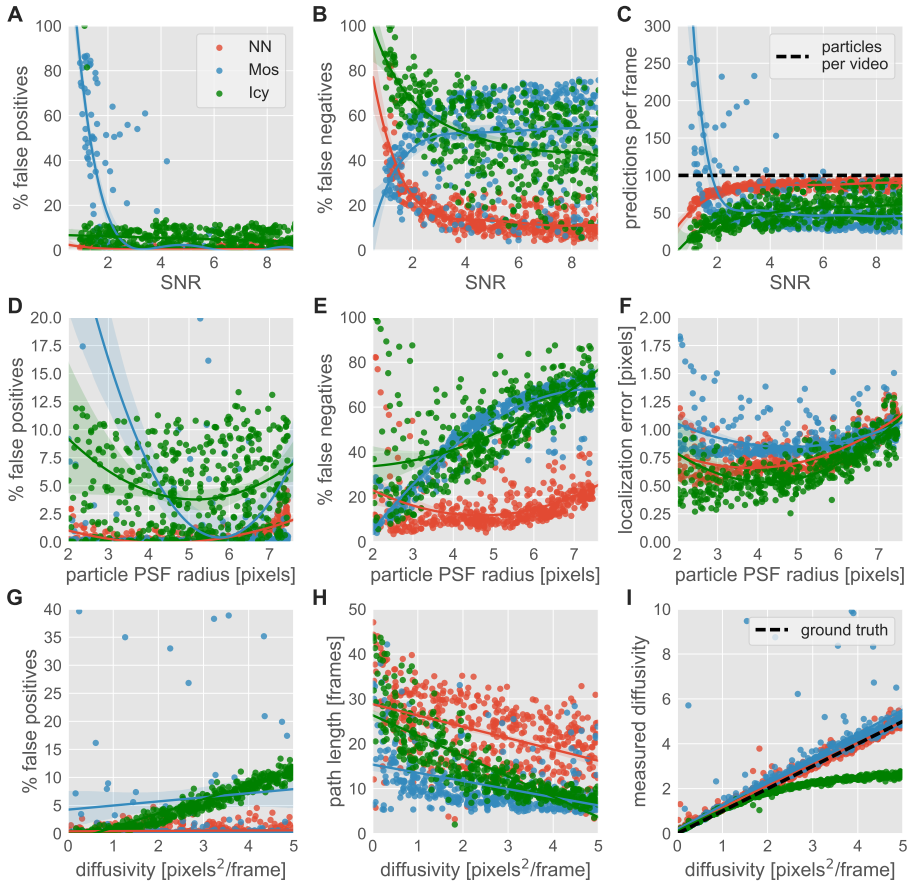

Figure 4: **Sensitivity analysis for randomized 2D synthetic test videos.** (A-C) 2D test results showing the (A) percentage of false positives, (B) percentage of false negatives and (C) predictions per frame vs. SNR. Mosaic shows a sharp rise in false positives for SNR < 2 (panel A), due to substantially more predictions than actual particles (panel C). Conversely, the neural net (NN) and Icy showed no increase in false positives at low SNR. (D-E) results showing the (D) percentage of false positives, (E) percentage of false negatives and (F) localization error vs. the PSF radius. (G-I) results showing the (G) percentage of false positives, (H) path duration, and (F) measured diffusivity vs. the ground truth particle diffusivity.

in excess of 50% across most SNR > 3 (Fig. 4B). In comparison, Icy possessed higher false positive rates than Mosaic at high SNR and lower false positive rates when SNR is decreased below 2.5, with a consistent ∼5% false positive rate across all SNR values (Fig. 4A). The false negative rates for Icy was greater than Mosaic at high SNR, and exceeded ∼40% for all SNR tested (Fig. 4B).

All three methods showed some minor sensitivity in the false positive rate and localization error to the PSF radius (Fig. 4D,F). (Note that the high sensitivity Mosaic displayed to changes in SNR made the trend for PSF radius difficult to discern.) Mosaic and Icy showed much higher sensitivity in the false negative rate to PSF radius, each extracting nearly 4-fold more particles as the PSF radius decreased from 8 to 2 pixels (Fig. 4E).

When estimating diffusivities, Icy exhibited increased false positive rates with faster moving particles (Fig. 4G), likely due to the linker compensating for errors made by detection algorithm. In other words, while the linker was able to correctly connect less-mobile particles without confusing them with nearby false detections, when the diffusivity rose, the particle displacements tended to be larger than the distance to the nearest false detection. This also caused Icy to underestimate the diffusivity (Fig. 4I) when $D > 2$.

In contrast to Mosaic and Icy, the neural network tracker possessed a far lower mean

false positive rate of ∼0.5% across all SNR values tested (Fig. 4A). The neural network tracker was able to achieve this level of accuracy while extracting a large number of paths, with <20% false negative rate for all SNR > 2.5 and only a modest increase in the false negative rate at lower SNR (Fig. 4B). Importantly, the neural network tracker performed well under low SNR conditions by making fewer predictions, and the number of predictions made per frame are generally in reasonable agreement with the theoretical maximum (Fig. 4C). Since the neural network was trained to recognize a wide range of PSFs, it also maintained excellent performance (<1% false positive, <20% false negative) across the range of PSF radius (Fig. 4E). The neural network tracker possessed comparably good localization error as Mosaic and Ivy, less than one pixel on average and never more than two pixels, even though true positives were allowed to be as far as 5 pixels apart (Fig. 4F). The neural network tracker was able to accurately track and reproduce the diffusivities across a wide range of particle speeds (Fig. 4G-I).

Performance on simulated 3D videos

When analyzing 3D videos, Mosaic and Icy were able to maintain roughly comparable false positive rates (∼5-8%) as analyzing 2D videos (Fig. 5A). Surprisingly, analyzing 3D videos with the neural network tracker resulted in an even lower false positive rate than 2D videos, with ∼0.2% false positives. All three methods capable of 3D tracking exhibited substantial improvements in reducing false negatives, reducing localization error, and increasing path duration (see Fig. 5B-D). Strikingly, the neural network was able to correctly identify an average of ∼95% of the simulated particles in a 3D video, i.e., <5% false negatives, with the lowest localization error as well as the longest average path duration among the three methods.

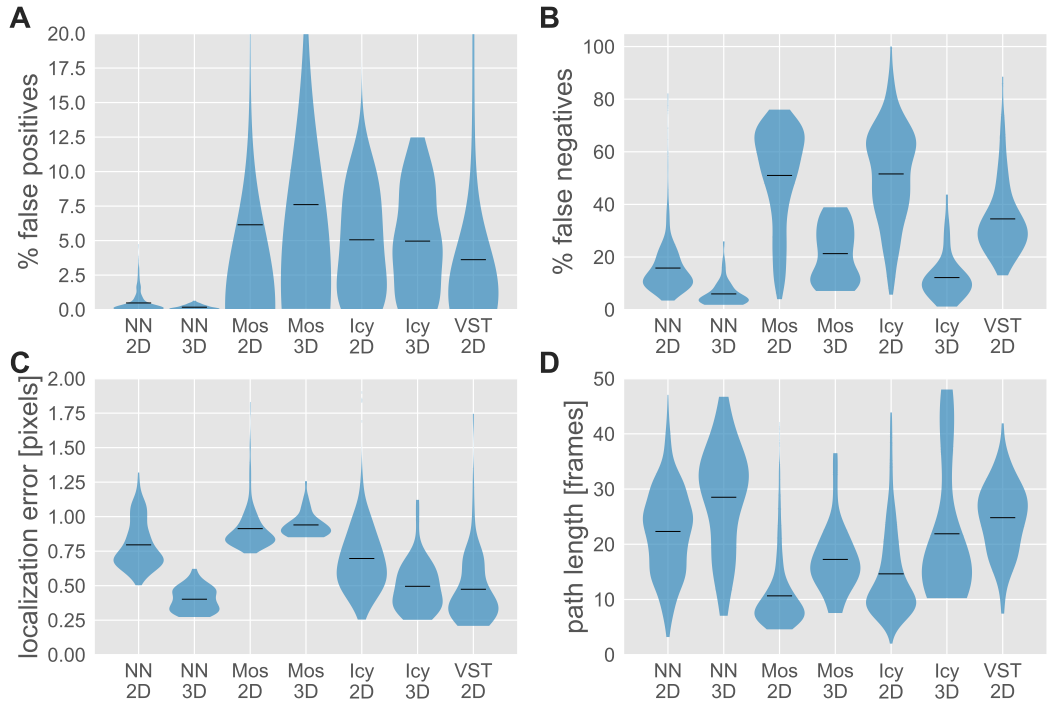

Figure 5: **Violin plots showing the performance on 3D test videos for each of the four methods: the neural network tracker (NN), Mosaic (Mos), Icy, and VST.** Performance with 2D simulated videos are included for comparison. The solid black lines show the mean, and the thickness of the filled regions show the shape of the histogram obtained from 500 (50) randomized 2D (3D) test videos. Note that the VST results only included 100 test videos.

Performance on experimental 2D videos

Finally, we sought to evaluate the performance and rigor of the neural network tracker on experimentally-derived rather than simulated videos, since the former can include spatiotemporal variations and features that might not be captured in simulated videos. Because analysis from the particle traces can directly influence interpretations of important biological phenomenon, the common practice is for the end-user to supervise and visually inspect all traces to eliminate false positives and minimize false negatives. Against such rigorously verified tracking, the neural net tracker was able to produce particle paths with comparable mean squared displacements across different time scales, alpha values, a low false positive rate, greater number of traces i.e. decrease in false negative, and comparable path length (see Fig. 6). Most importantly, these videos were processed in less than one twentieth of the time it took to manually verify them, generally taking 30-60 seconds to process a video compared to 10-20 minutes to verify accuracy.

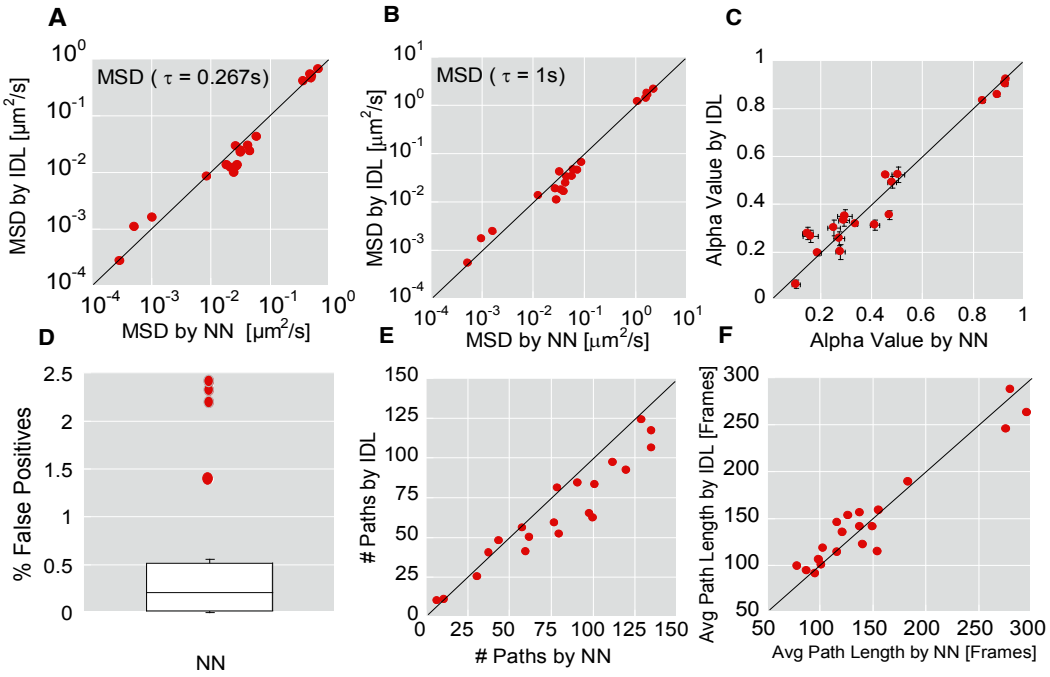

Figure 6: **Comparison of human tracked (assisted by the commercially available software package IDL) and neural network tracked output.** Ensemble-averaged geometric mean square displacements ($\langle$MSD$\rangle$) at a time scale (τ) of (A) 0.267s and (B) 1s. (C) alpha value (D) percentage of false positives normalized by path-length (E) number of particles tracked (F) average path duration per particle. The error bars in (C) represent standard error of the mean. The box plot in (D) shows symbols for the outliers above the 80th percentile of observations. The data set includes 20 different movies encompassing muco-inert 200 nm PEGylated polystyrene beads 200 nm carboxylated beads, HIV virus-like particles and herpes simplex virus in human cervicovaginal mucous.

Discussion

Although tracking the motion of large, bright, micron-sized beads is straightforward, it remains exceptionally difficult to rapidly and accurately obtain traces of entities, such as ultrafine nanoparticles and viruses, that are sub-micron in size. Sub-micron particles can readily diffuse in and out of the plane of focus, possess low SNR or significant spatial heterogeneity, and undergo appreciable photo-bleaching over the timescale of imaging. Accurate conversion of videos to particle paths for these entities necessitates extensive human intervention; it is not surprising to spend 10-20x more time on extracting path data from videos

than the actual video acquisition time. Worse, substantial user variations is common even when using the same software to analyze the same videos (Fig. 2). Analysis throughput is further limited by the fact that students/users can rarely process more than 10-20 videos per day without encountering 'tracker fatigue' that hinders their decision making. These challenges have strongly limited particle tracking, preventing it from becoming a widely-used tool in physical and life sciences.

To tackle these challenges, we developed here a CNN comprised of over 50,000 parameters, and employed deep learning to optimize the network against a diverse array of video conditions. The end product is a fully automated tracker can consistently analyze both 2D and 3D videos with a remarkably low false positive rate, and lower false negative rate, lower localization error and longer average path lengths than a number of the leading particle tracking software. Importantly, because the tracker requires no user input, the output is absolute and reproducible, which not only improves the rigor and reproducibility of all particle tracking experiments within a lab but also enable different labs across the world to compare particle tracking data in an objective manner. Finally, the neural network tracker greatly increases the throughput of converting videos into particle position time series, which addresses in our opinion the biggest bottleneck limiting the applications of particle tracking.

CNN has become the state-of-the-art for object recognition in computer vision, outperforming other methods for many imaging tasks. A number of research groups are beginning to apply machine learning to object identification and tracking, primarily involving 'hand crafted' features that in essence serve as a set of filter banks for making statistical measurements of an image, such as mean intensity, standard deviation and cross correlation. These features are used as inputs for a support vector machine, which is then trained using machine learning. The use of hand-crafted features substantially reduces the number of parameters that must be trained. In contrast, we have developed our network to be trained end-to-end, or pixels-to-pixels, so that the input is the raw imaging data, and the output is a probabilistic classification of particle versus background at every pixel, with all parameters optimized via machine learning. Importantly, we have designed our network to effectively perform convolutions in both space and time (past, present and future) to infer particle locations.

Currently, nearly all particle tracking is performed on 2D videos, despite the fact that 2D PT generally can only capture the vast majority of highly mobile particle species for exceedingly short duration. It is dangerous, and frequently incorrect, to extrapolate measurements at short timescales to make predictions of behaviors spanning long time scales. Indeed, the small number of observations of mobile species with 2D PT preclude the use of sophisticated Bayesian and Maximum Likelihood Estimators to perform model selection and get tight bounds on parameters of best-fit models, thus losing the power of prediction beyond experimental timescales. Naturally, 3D imaging can substantially improve particle tracking quality, and provide critical insights into numerous physiological phenomena that cannot be adequately captured by 2D PT. Nevertheless, visualizing 3D volumetric time series data is a significant challenge for the human eye, which is limited to perceiving 2D projections, and 3D videos contain at least 10-50 times more data than a comparable 2D video. Although software-assisted tracking is widely available for 3D videos, the excessive time needed to verify accurate tracking, coupled with data storage requirements, have prohibited most research groups from attempting 3D particle tracking, even though the technical capability to do so has existed for decades. The few groups that acquire videos in 3D generally only do so to choose the most visually optimal plane for 2D tracking analysis. By requiring no user-input, we believe the neural network tracker can tackle the longstanding challenge of analyzing 3D videos, and in the process encourage broad adoption of 3D PT for applications previously deemed inaccessible. Likewise, by enabling rapid and user-independent analysis, the neural network tracker should in theory also enable fully automated particle tracking experiments.

Acknowledgements

Financial support was provided by The David and Lucile Packard Foundation (2013-39274, S.K.L.), NSF CAREER Award (DMR-1151477, S.K.L.), and funding from the Eshelman Institute of Innovation. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

[1] Y.-Y. Wang, A. Kannan, K. L. Nunn, M. A. Murphy, D. B. Subramani, T. Moench, R. Cone, and S. K. Lai, "Igg in cervicovaginal mucus traps hsv and prevents vaginal herpes infections," *Mucosal immunology*, vol. 7, no. 5, pp. 1036–1044, 2014.

[2] K. L. Nunn, Y.-Y. Wang, D. Harit, M. S. Humphrys, B. Ma, R. Cone, J. Ravel, and S. K. Lai, "Enhanced trapping of hiv-1 by human cervicovaginal mucus is associated with lactobacillus crispatus-dominant microbiota," *MBio*, vol. 6, no. 5, pp. e01 084–15, 2015.

[3] Y. Wang, D. Harit, D. Subramani, H. Arora, P. Kumar, and S. Lai, "Influenza-binding antibodies immobilize influenza viruses in fresh human airway mucus," *European Respiratory Journal*, vol. (in press), 2016.

[4] S. K. Lai, D. E. O'Hanlon, S. Harrold, S. T. Man, Y.-Y. Wang, R. Cone, and J. Hanes, "Rapid transport of large polymeric nanoparticles in fresh undiluted human mucus," *Proceedings of the National Academy of Sciences*, vol. 104, no. 5, pp. 1482–1487, 2007.

[5] M. Yang, S. K. Lai, Y.-Y. Wang, W. Zhong, C. Happe, M. Zhang, J. Fu, and J. Hanes, "Biodegradable nanoparticles composed entirely of safe materials that rapidly penetrate human mucus," *Angewandte Chemie International Edition*, vol. 50, no. 11, pp. 2597–2600, 2011.

[6] P. A. Vasquez, C. Hult, D. Adalsteinsson, J. Lawrimore, M. G. Forest, and K. Bloom, "Entropy gives rise to topologically associating domains," *Nucleic acids research*, p. gkw510, 2016.

[7] T. Mason, K. Ganesan, J. Van Zanten, D. Wirtz, and S. C. Kuo, "Particle tracking microrheology of complex fluids," *Physical Review Letters*, vol. 79, no. 17, p. 3282, 1997.

[8] D. Wirtz, "Particle-tracking microrheology of living cells: principles and applications," *Annual review of biophysics*, vol. 38, pp. 301–326, 2009.

[9] D. Chen, E. Weeks, J. C. Crocker, M. Islam, R. Verma, J. Gruber, A. Levine, T. C. Lubensky, and A. Yodh, "Rheological microscopy: local mechanical properties from microrheology," *Physical review letters*, vol. 90, no. 10, p. 108301, 2003.

[10] I. Wong, M. Gardel, D. Reichman, E. R. Weeks, M. Valentine, A. Bausch, and D. Weitz, "Anomalous diffusion probes microstructure dynamics of entangled f-actin networks," *Physical review letters*, vol. 92, no. 17, p. 178101, 2004.

[11] T. A. Waigh, "Microrheology of complex fluids," *Reports on Progress in Physics*, vol. 68, no. 3, p. 685, 2005.

[12] N. Flores-Rodriguez, S. S. Rogers, D. A. Kenwright, T. A. Waigh, P. G. Woodman, and V. J. Allan, "Roles of dynein and dynactin in early endosome dynamics revealed using automated tracking and global analysis," *PloS one*, vol. 6, no. 9, p. e24479, 2011.

[13] K. M. Schultz and E. M. Furst, "Microrheology of biomaterial hydrogelators," *Soft Matter*, vol. 8, no. 23, pp. 6198–6205, 2012.

[14] L. L. Josephson, E. M. Furst, and W. J. Galush, "Particle tracking microrheology of protein solutions," *Journal of Rheology*, vol. 60, no. 4, pp. 531–540, 2016.

[15] M. T. Valentine, P. D. Kaplan, D. Thota, J. C. Crocker, T. Gisler, R. K. PrudâĂŹhomme, M. Beck, and D. A. Weitz, "Investigating the microenvironments of inhomogeneous soft materials with multiple particle tracking," *Physical Review E*, vol. 64, no. 6, p. 061506, 2001.

[16] S. K. Lai, Y.-Y. Wang, K. Hida, R. Cone, and J. Hanes, "Nanoparticles reveal that human cervicovaginal mucus is riddled with pores larger than viruses," *Proceedings of the National Academy of Sciences*, vol. 107, no. 2, pp. 598–603, 2010.

[17] S. K. Lai, Y.-Y. Wang, D. Wirtz, and J. Hanes, "Micro-and macrorheology of mucus," *Advanced drug delivery reviews*, vol. 61, no. 2, pp. 86–100, 2009.

[18] J. C. Crocker and D. G. Grier, "Methods of digital video microscopy for colloidal studies," *Journal of colloid and interface science*, vol. 179, no. 1, pp. 298–310, 1996.

[19] N. Chenouard, I. Smal, F. De Chaumont, M. Maška, I. F. Sbalzarini, Y. Gong, J. Cardinale, C. Carthel, S. Coraluppi, M. Winter *et al.*, "Objective comparison of particle tracking methods," *Nature methods*, vol. 11, no. 3, p. 281, 2014.

[20] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[21] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.

[22] T. Savin and P. S. Doyle, "Static and dynamic errors in particle tracking microrheology," *Biophysical journal*, vol. 88, no. 1, pp. 623–638, 2005.

[23] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks." in *Aistats*, vol. 15, no. 106, 2011, p. 275.

[24] R. Lumia, "A new three-dimensional connected components algorithm," *Computer Vision, Graphics, and Image Processing*, vol. 23, no. 2, pp. 207–217, 1983.

[25] R. Parthasarathy, "Rapid, accurate particle tracking by calculation of radial symmetry centers," *Nature Methods*, vol. 9, no. 7, pp. 724–726, 2012.

[26] K. Jaqaman, D. Loerke, M. Mettlen, H. Kuwata, S. Grinstein, S. L. Schmid, and G. Danuser, "Robust single-particle tracking in live-cell time-lapse sequences," *Nature methods*, vol. 5, no. 8, pp. 695–702, 2008.

[27] X. Xiao, V. F. Geyer, H. Bowne-Anderson, J. Howard, and I. F. Sbalzarini, "Automatic optimal filament segmentation with sub-pixel accuracy using generalized linear models and b-spline level-sets," *Medical image analysis*, vol. 32, pp. 157–172, 2016.

[28] J.-C. Olivo-Marin, "Extraction of spots in biological images using multiscale products," *Pattern recognition*, vol. 35, no. 9, pp. 1989–1996, 2002.

[29] N. Chenouard, I. Bloch, and J.-C. Olivo-Marin, "Multiple hypothesis tracking for cluttered biological image sequences," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 11, pp. 2736–3750, 2013.

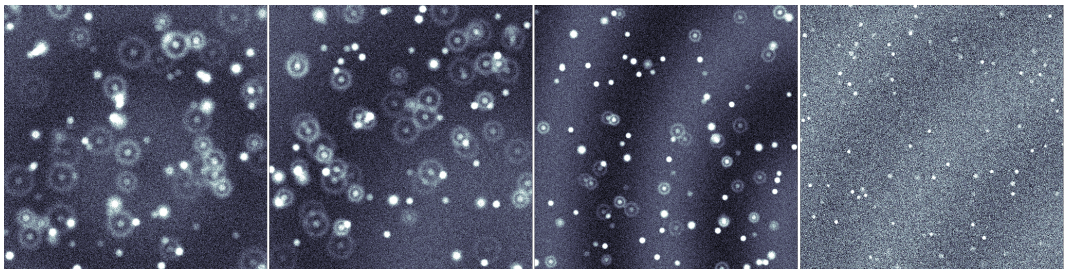

Figure 7: Supplemental: Sample frames from four different synthetic test videos.

Supplemental: Parameter values for tracking software used in the synthetic video tests

Sample frames of the synthetic test videos can be seen in Fig. 7. All of the 2D and 3D videos were tracked using the same set of parameter values. Through experimentation, these parameter values showed the best performance overall. The neural net tracker uses one parameter in its linking method (for collecting particle localizations into paths). The standard deviation for particle displacements was set to $\sigma = 20$. For Mosaic, we used a custom ImageJ macro, using three parameter values, to batch process the test videos. The particle detection parameters were

$$\text{radius} = 8, \quad \text{cutoff} = 0, \quad \text{percentile} = 0.8$$

For ICY, we used a custom javascript script for batch processing, which only required parameter values for its partical localization method. The particle detection parameters were

$$\text{scale}_1 = 0, \quad \text{scale}_2 = 0, \quad \text{scale}_3 = 50, \quad \text{scale}_4 = 100$$

For linking, we specified that particles with PSF radius < 2 (the minimum size in the test videos) be filtered, and that the ICY linker should assume all particles move by standard Brownian motion.